[SQUEAKING]

[RUSTLING]

[CLICKING]

JONATHAN GRUBER: OK, so we are going to dive in today and talk about where demand curves come from. The motivation for the last couple of lectures was that we showed you a demand curve and a supply curve, and I promised I'd tell you what the basis for those demand supply curves are.

The last two lectures we started building that basis. And today we'll show you where demand curves actually come from. So let's actually derive-- let's start by actually deriving a demand curve. OK. So deriving a demand curve. OK.

So basically, demand curve, as we said in the first lecture, is the relationship between from the price charged for a good and the consumer's desired quantity of the good. So what we'll do is return to the example before and show you where that comes from.

So remember the example before. Remember, the utility function, which was the square root of s times c, slices times cookies. We had a budget constraint, which was y is equal to the price of slices times the number of slices, plus the price of cookies times the number of cookies. And we assumed to start that y was 24, the price of slices was 4, and the price of cookies was 2.

So how do we derive a demand curve? Step 1 is we solve for optimal consumption. Solve for s star c star. Solve for the optimal consumption pair. How do we do that? Well, as we said last time, constrained optimization or the point where you can-- the highest indifference curve you can reach given your budget constraint leads you to set MRS equal to MRT.

That's the tangency condition. That's the point where you're as well off as possible. Remember, it's marginal thinking in economics. It's the next unit. You want to consume pizza and cookies until the next unit makes you as happy as the price the market charges.

How happy does it make you? Well, that's your marginal rate of substitution. That's minus s over c. What does the market charge? Well, that's minus the price ratio, which is minus 1/2. OK? So basically, we derived that last time. Yes, we derived that two times ago. That is the constrained optimization.

Now we then-- so last time. But we go to the next step now, which is we say, well, what does that imply for optimal slices and cookies? Well, we just plugged that into the budget constraint. We have this equation, which is one equation in two unknowns. We also have this equation, which is that 24 equals 4s plus 2c.

So now we have two equations and two unknowns. OK. We know how to do that. OK. And we solve that, we get that c star equals 6, and s star equals 3. We just solve those two equations and two unknowns. And then we get c star is 6 and s star is 3. OK?

So that is illustrated in the top panel of figure 1. In the future, let's not make it so compressed. I'm OK with killing a few more trees to make sure our fine students don't have to hurt their vision.

So basically that is, if you look at in figure 4.1 at the very top, you look at budget constraint one, that is our original budget constraint we just had. And the tangency, the optimum point is at point A. You want six cookies and three slices of pizza. This is reviewing what we did before. Questions about that though? Because this is some math you're going to do over and over and over again until you're sick of it. Questions about that?

OK. So now let's ask, what happens if the price of cookies changes? What happens if the price of cookies rises to 3? So let's say that the price of cookies goes from 2 to 3. Well, the new slope of the budget constraint is now instead of minus 1/2, it's minus 1/3.

So your optimization problem is set minus s over c equal to minus 1/3, not minus 1/2. Your budget constraint is now 24 equals 4s plus 3c. Once you have two equations and two unknowns, you've solved that. And you get that c star equals 4, s star equals 3. Yeah, question. Speak up so everyone can hear you.

AUDIENCE: Isn't it negative?

JONATHAN GRUBER: Oh, negative s over c. I'm sorry, what is your question?

AUDIENCE: Isn't it negative 3 over 4?

JONATHAN GRUBER: Is it negative? Did I do that wrong? Yes, you're right. My bad. Negative 3 over 4. I should never do that when I look at my notes. Good catch. Negative 3 over 4. OK? So you solve that and you end up getting the optimum number of cookies is 4, optimum price of slice of pizza is 3.

That is point B on the top diagram. How did we get there? Well, mathematically, I showed you. Subject to one mathematical mistake. I showed you how to get there. Graphically, how do we get there? Well, we draw the new budget constraint, BC2. And we find the tangency between the indifference curve and that new budget constraint.

BC2 offers a smaller opportunity set, as we discussed last time. So you'll end up with fewer goods. You have a smaller opportunity set, and you do. You end up with the same amount of pizza but with fewer cookies.

Likewise, what if the price of cookies fell to 4/3? What if the price of cookies went from 2 to 4/3? Well, once again, you can do the math, and you'll find that basically when you do that, you're going to want three slices and nine cookies.

And graphically, you can see that that's a outward shift in the budget constraint to BC3. That's the rightmost budget constraint. Your indifference curve is tangent with that budget constraint at nine cookies and three slices of pizza.

So what's happening here? We are shifting the price of cookies and showing you, based on tools you've learned, how that impacts how many cookies you want and how many slices you want. In other words, we're drawing the demand curve.

Look at the translation from the top figure to the bottom figure. The bottom figure has on it cookies on the x-axis and the price of cookies on the y-axis. Well, that's a demand curve. We've just drawn it. I've just shown you for three different price of cookies, how many different cookies you want.

So we have now derived the demand curve. We derive the demand curve by optimizing your utility at different prices and finding which quantity you want at each of those different prices. And that's the demand curve. It's the quantity you want at each price. So that's where the demand curve comes from. OK. Yeah?

AUDIENCE: Something that I noticed is that when we vary the price of cookies, the number of slices of pizza that we want is the same.

JONATHAN GRUBER: Great point. Someone noticed. And that as we varied the price of cookies, the number of slices of pizza did not change. That is not a general law. That is a specific feature of this utility function. In general, when you change one price, both goods will change.

What's nice-- one reason we use this convenient utility function is it has some nice mathematical properties, and one is it makes life easy by having the slice of pizza stay the same. So essentially what we're saying here is that we have four pizza, we have what we call a flat price consumption curve.

OK? Which is based on-- a cross-price consumption curve, which is basically when the price of one good changes, the demand for the other good doesn't change. And that's just a feature of this particular utility function. But it's not general. Now in fact, you can-- yeah, go ahead.

AUDIENCE: I find a negative answer that's usually the marginal utility of the cookies.

JONATHAN GRUBER: Exactly. So I skipped a step, which is defining the MRS. I defined that last time, or two times. I don't remember. Two times ago, I think. The MRS is, remember, the negative of the ratio of the marginal utility of cookies to the marginal utility of pizza. And that's minus.

So the way we got this, we did two lectures ago, we basically took the derivative of the utility function first with respect to cookies then with respect to pizza. We got the two different marginal utilities. And then we took the ratio that gives the marginal rate of substitution.

All right? Good question. Good to clarify where that came from. So this is building on the work we did two lectures ago. Other questions? OK. So that is where demand curves come from. Fun exercise at home, do the same exercise, changing the price of pizza. Do the reverse exercise. You will find as well there's a flat-- there's no effect to the number of cookies. But you will find an effect on the number of slice of pizza. And you can show how that will vary.

OK. Now that we know what demand curves are, I want to talk about what do they look like. Let's talk about the shape of demand curves. What drives what they look like? OK. And what drives what demand curves look like is going to be what we call the elasticity of demand.

The elasticity of demand, epsilon, is going to be the percentage change in quantity, delta q over q, per percentage change in price, or in a continuous time, dq dp. And basically, we're going to generally work in discrete space, but the math you can do continuously.

I write it both ways because you're going to-- the examples we'll use will do in discrete space. It's a little easier to see it. But basically, it's the derivative of the quantum with respect to the price. It's the percentage change in quantity with respect to price. So approximately equals the derivative.

So basically, that is the elasticity of demand. This tells you how responsive people's demand is to the price they have to pay. Once again, I find it very useful to fix ideas in this class to think about extreme cases that never existed in reality, but can help you understand the intuition.

So let's look at figure 4.2. Figure 4.2 would be a case of perfectly inelastic demand. That would be eta equals 0. That we call perfectly inelastic demand. That's a vertical demand curve.

In this case, demand does not respond to prices. Can anyone think of a good that even could come close to a perfectly inelastic demand or feature, even features of a good. Yeah.

AUDIENCE: Like gas.

JONATHAN GRUBER: Like what?

AUDIENCE: Gas.

JONATHAN GRUBER: Gas? Gas would be-- well, it's-- you're right. In the sense-- it's a very interesting example. In the sense, it's inelastic because you got to get to work. On the other hand, there's a bus, there's your bike. You could walk, you could buy a new car. So it's not perfectly inelastic. But it's somewhat. What else? Yeah.

AUDIENCE: Water.

JONATHAN GRUBER: Water would be something. But once again, not perfectly because you can drink Coke, you can drink juice, you can drink other things. You're on the right track. Yeah.

AUDIENCE: Insulin.

JONATHAN GRUBER: Insulin is the classic example we use. There's no substitute. When you're having diabetic-- when you're diabetic trauma for insulin, you will not say, it's $1 more, I'm just going to die. OK? We think it's fairly-- it's not perfectly-- nothing's perfectly inelastic. At some point if insulin was $1 million, you might just not buy it because you can't afford it. But it's certainly locally inelastic. That'd be an example of--

But once again, these are great examples of things which-- the examples given are goods which should be somewhat inelastic but not perfectly. Likewise, in figure 4.3, we consider a good with perfectly elastic demand, which is epsilon equals negative infinity.

Perfectly elastic demand. That is, any infinitesimal change in price causes an infinite change in quantity. Once again, doesn't exist in reality. But what types of goods would have very, very, very elastic demand? Yeah.

AUDIENCE: Luxury goods.

JONATHAN GRUBER: Luxury goods would have elastic demand. Because basically, you don't really need them. What else? Yeah.

AUDIENCE: [INAUDIBLE] brands.

JONATHAN GRUBER: Yeah. So McDonald's versus Burger King. Now, we may have our preferences, but the bottom line is that if McDonald's charge twice as much as Burger King for a burger, everybody would just go to Burger King. Basically there's a point at which there's so-- and essentially, what is-- oh, question?

AUDIENCE: [INAUDIBLE]

JONATHAN GRUBER: Oh, no, we're getting examples. But here's the key feature that differentiates elasticity. It's substitutability. Goods with no substitutes are not elastic. Goods with lots of substitutes are very elastic. That's why I immediately respond to gas by talking about all-- that's so embedded in my brain, that when someone brought up gas, what did I go to? I went to all the substitutes. Because more substitutes there are, the more elastic the good is.

So basically, that's essentially how to think about elasticity. Think about substitutability. The degree of substitutability is the degree of elasticity. And obviously, everything in the world is going to lie between these two extremes.

There's nothing in the world that's ever purely inelastic and purely elastic. But things are to lie between these two extremes, and where they lie will be a function of how substitutable they are. Questions about that?

Now, let me make a point that we are going to in this class once again be cheating a little bit about linear versus nonlinear demand curves. A constant elasticity demand curve would have to be nonlinear. If you don't see that right away, it's something you should go home and think about. A constant elasticity--

But we're going to talk about curves. We're going to talk about one elasticity and draw linear curves. So once again, just think of them as very local approximations. So it's approximately right. So don't get too thrown off by the nonlinear/linear part of this stuff. And if one thing if it's unclear, ask. Yeah.

AUDIENCE: For elastic equation for continuous states, would that be ap or qp?

JONATHAN GRUBER: No. It's the elasticity, it's how much incremental change in price changes quantity. So now let's talk about another topic, which is we've talked-- remember we talked about what moves budget constraints. We talked about prices changing, income changing.

So let's talk about what happens to demand curves when income changes. Income and demand curves. So now let's go to figure 4.4. Once again, we start at point A. That is our original setup with the budget constraint 24 equals 4s plus 2c, and the utility function equals square root of c. And that gets the solution at point A.

Now, imagine that your income rises from $24 to $32. Nothing else has changed. Your preference is the same, price is the same. But your income has risen from $24 to $32. Well, in that case, the optimization is now you're still setting minus s over c to a price ratio.

You're still setting minus s over c to the price ratio minus 1/2. That hasn't changed because the price ratio hasn't changed and your preferences haven't changed. What's changed is the second equation, which you've already illustrated is now 32 equals 4s plus 2c.

So the tangency condition hasn't changed. What's changed and really goes up is the budget constraint. And we solve for that. You'll now see you get that c star equals 8, and s star equals 4. So you change your consumption but not because prices have changed. You're richer. You want more. That is exactly what we see at point B.

Point B shows-- remember, when your income goes up, there's a parallel shift outward in the budget constraint. Income goes up, parallel shift out in the budget constraint. You still want to find the tangency. That tangency is still going to solve this condition because prices and preferences haven't changed. It's just going to be a higher quantity. And that's what happens at point B. Same tangency at a higher quantity.

Likewise, at point C, we show what happens if your income falls to $16. Then you want less pizza and fewer-- less pizza and fewer cookies. Four cookies and two slices. What does that lead us to? That leads us to a relationship now between your income and the quantity demanded. And that relationship is shown in the second panel of figure 4.4. We call this the Engel curve.

The Engel curve, which is the name for the relationship between income and quantity demanded. And you can draw an Engel curve for s and Engel curve for c. Here we're drawing the Engel curve for c.

And the Engel curve, the slope of the Engel curve is what we call the income elasticity of demand, or gamma, which is delta q over q, over delta y over y, or approximately dq dy. OK.

So basically, this is instead of how does demand change with price, which is the demand curve, this is, how does demand change with income. Or what we call the Engel curve. And we're going to come in a few minutes about why this is important. But let's talk about the properties of gamma.

Gamma is much more interesting than epsilon. Epsilon, some negative number, 0 to minus infinity. Gamma can have a much larger range. In fact, gamma is usually greater than 0. In fact, when gamma is greater than 0. We call them normal goods. Normal goods are goods where gamma is greater than 0.

But gamma does not have to be greater than 0. Well, epsilon has to be less than or equal to 0. Gamma does not. Gamma could actually be less than 0. Does anyone know how that could happen? How could it happen that you get richer and want less of something? What's an example? Yeah. It's what's called an inferior good. OK, but give an example.

AUDIENCE: [INAUDIBLE]

JONATHAN GRUBER: Right. Right. Exactly. So as you get richer, you want to get rid of things that you're only eating because they were cheap. What's an example of that? Yeah.

AUDIENCE: Ramen.

JONATHAN GRUBER: Ramen is the example of that. Ramen is the ultimate inferior good. No one wants to eat ramen. I realize this fancy ramen restaurants, but nobody should go there. OK? People want to eat ramen-- people eat ramen because it's cheap.

No, in all seriousness, it is amazing that ramen has become a restaurant food after all. It is in other countries, I realize. But most consumption of ramen is done by cheap packets of the supermarket. As you get richer, you consume less ramen. You might not consume zero. You might still like it, but you're not eating it five days a week like when you're starving entrepreneur trying to make your own business.

So ramen is an example of an inferior good, a good that as you get richer, you have less of it because basically you're only eating it because it was cheap. Now even within normal goods, we have a distinction, which is within normal goods, we distinguish between two things, which is gamma greater than 1, and gamma less than 1.

What do you think we call goods with gamma greater than 1 versus gamma less than 1? And this relates to examples that were given before. Yeah.

AUDIENCE: Luxury goods.

JONATHAN GRUBER: Are which one?

AUDIENCE: Greater than 1.

JONATHAN GRUBER: Greater than 1. So basically, luxury goods are where as your income rises, the percent of your income spent on the good goes up. Let's be very clear. As your income rises, you consume more of everything generally.

But it's, how does the budget share change? Do you spend more and more of your budget on that good as you get richer? And that's what we call a luxury good. So a classic luxury good would be eating at restaurants.

When you're poor, you have to buy-- you have to make your own dinners because the restaurant's too expensive. As you get richer, a larger and larger share of your budget will be consumed by eating at restaurants. That's a luxury good. What about gamma less than 1? What we call that? Got any ideas? Gamma less than 1. Yeah.

AUDIENCE: Giffen good.

JONATHAN GRUBER: What's that?

AUDIENCE: Giffen good.

JONATHAN GRUBER: No. Giffen good is-- well, we'll come back to that. You're one step ahead. Gamma less than 1. Simpler than that. So if these are luxuries, what are these? Necessities. Sorry for the small writing. Necessities. These are goods where as your income goes up, you don't have less of them.

Ramen is an inferior good. You literally have less ramen. But this is like eating at home. Eating at home, you don't necessarily-- you don't spend less on eating at home as you get richer, but you might spend a lower share. Or this could be clothes. Or this could be rent.

Things where if I get rich, you get a nicer apartment. But if you get really rich-- when you're poor, your apartment is 60% of your income. When you're rich, it's going to fall. Yeah.

AUDIENCE: So for necessities, it's less than 1, but more than 0, right?

JONATHAN GRUBER: More than zero. Exactly. Exactly. It's less than 1 but more than zero. That's a great point. These are both greater than 1. This is less than 1, but greater than zero. If it's less than 0, it's inferior.

So that's the interesting distinction between necessity and inferior. Let me explain. It's very clear. Inferior is-- it's very unclear, I mean. Inferior is as your income goes up, you literally have less of it. Necessity as your income goes up, you have more of it. But it goes-- but it rises slower than your income.

So rent is the perfect example. Richer people have nicer apartments than poorer people, but richer people spend a smaller share of their budget on apartments than do poorer people. It's a necessity. Richer people have nicer watches than poor people, and it's a rising share of their budget as they get richer because they don't have a nice watch when they're poorer. And as they get rich, they buy nicer, nicer watches. Yeah.

AUDIENCE: Is there a way to connect [INAUDIBLE]?

JONATHAN GRUBER: We're going to do that in one minute. Other questions? These are the primitives that allow us to go on to do the fundamental goal of this lecture, which is to get behind the mechanics of what happens when you change prices.

Does everyone get this? Can I put this up? OK. What happens? What's the mechanics of what happens when you change prices? Now you said-- you say, John, you just showed us what happens when you change prices. The price goes up, people want less of it. The price goes down, people want more of it. You derived the demand curve, Let us go a half hour early.

And I say, no, I will not let you go a half hour early, because there's actually a very important underlying mechanics going on when the price changes. And this mechanics is largely not so relevant for most goods we do, but for some goods if we delay the semester, it'll be critical. And that's why it's important to understand the mechanics.

And this is basically the fact that when a price changes, when a price changes-- so this is called mechanics of a price change. When a price changes, two things happen simultaneously. You may not realize it, but actually when a price changes, two things happen.

The first is what we call the substitution effect. The substitution effect, which is when relative prices change, you want to shift away from the good that's gotten relatively expensive towards the good that's gotten less expensive. That's the substitution effect. And you say yeah, that makes sense. That's a price change.

But the second one is the income effect, which is when a price changes, you're also effectively richer or poorer. And that also affects your demand. So when a price changes, two things happen. Goods become differentially attractive and you're richer or poorer. And both things drive your response to that price change. And it turns out, understanding the difference between them will end up being critical for a number of things we'll do this semester. And that's why I'm bothering telling you so.

So to understand this, we're going to go through one of the more complicated examples of this class. So that's why I'm leaving time for it. We're going to talk about decomposing a price change into an income and substitution effect.

And then we're going to do the math of this in section. Or did we-- yeah, and we'll do the mathematics of this in section. It's hard enough I'll do the graphics here in the intuition. The math will be covered in section.

So now let's go back to our classic example. We start at point A. We're on original budget constraint 1. So this is figure 4.5. Look at BC1. That's original budget constraint, and our indifference curve is tangent at point A. So as we've said a number of times this semester-- this lecture, we choose six cookies and three pizzas. That's our base case.

Now, imagine that the price of cookies rises from $2 to $3. Well, we know what happens. We solved for that already earlier in this lecture. When the price of cookies rise from $2 to $3, we know you move from point A to point C. You move from consuming six cookies and three slices of pizza to four cookies and three slices of pizza. We already solved for that.

Great. You already know that. Why aren't we done? We aren't done because that's composed of two effects. A substitution effect and income effect. How do we measure these? Well, the way we do it graphically is we say the substitution effect is technically defined as the change in the quantity purchased holding utility constant.

So it's delta q at u bar holding your utility constant. How would you change your mix of goods? How do we measure that? We measure by drawing an imaginary budget constraint, BC prime. This isn't a real budget constraint, it's imaginary. And we find the point.

And what is BC prime? BC prime has two features. First of all, it is parallel to the new budget constraint. It's parallel to the-- it's got the new price ratios. The slope of the new budget constraint of BC prime is the same as BC2.

So minus 3/4, right? That was the-- the slope was, yeah, minus 3/4. So the new budget constraint has a slope of minus 3/4. But it is tangent to the original indifference curve, the one that contains point A.

So once again, how do we measure the substitution effect? We measure it by drawing a new budget constraint, an imaginary budget constraint, that is parallel to the new budget constraint. It's the new price ratio, minus 3/4 instead of minus 1/2, but tangent to the old indifference curve.

And think about what that's doing. We're showing you the change in quantity you get at a constant utility. How do we keep utility constant? We stay in the same indifference curve. So by staying the same indifference curve, we're holding utility constant.

And we're showing you that even with utility constant, you would want fewer cookies. Why would you want-- why would you want fewer cookies? You'd move from six cookies to 4.89 cookies. Why? Because cookies have gotten more expensive relative to pizza. That's the pure substitution. It doesn't matter if you're rich or poor. It's just cookies have gotten more expensive, so you want fewer of them.

It's easy to understand this if we go directly to the second effect. If we go directly to the second effect, which is-- oh, well, no. Actually, let me do one thing first. I want to first prove to you that the substitution effect is always negative. The substitution effect is always negative.

So to prove that-- so graphically, we can prove that by just noting that if you're on the same indifference curve and you're tangent to a steeper budget constraint, you must be to the left. If I'm on the same indifference curve, the budget is steep, then I must be to the left of the original point. So it's got to be a negative effect.

Mathematically, think of this in five steps. Step 1, the substitution effect puts you at the tangency of the indifference curve and prices. We know at that tangency that MUc over MUs equals pc over ps. I just canceled the negatives.

We know those are equal. At any tangency, we know those are equal. At every tangency. We know that this has gone up. Price of cookies has gone up relative to pizza slices. Right? Once again, jump in if anything's not clear.

If that goes up, what has to happen? This has to go up. They have to be equal. The right side went up, so the left side has to go up. How do you make the marginal utility ratio go up? How do you make it so that the ratio of marginal cookies to marginal slices goes up? What do you do? How can you operationalize that? Well, how do you make the numerator higher. How do you make the marginal utility of cookies bigger? Yeah.

AUDIENCE: By marginal utility.

JONATHAN GRUBER: Yeah. Remember, marginal utilities are negative functions of quantity. I said they're confusing. The fewer cookies you have, the more each cookie is worth. If you have fewer cookies, that means you have to have more slices. If you have more slices, what does that do to the marginal utility of slices? It lowers it.

So by shifting from cookies to slices, you lower the ratio-- I'm sorry, you raise the ratio of marginal utilities. You raise the numerator and lower the denominator through that shift. So you therefore say that if you're going to be tangent and you're going to have a higher price ratio, you must have fewer cookies and more slices by definition. And that's the proof that substitution effects are negative. OK? Questions about that?

AUDIENCE: So the substitution effect is if the price goes up?

JONATHAN GRUBER: The price goes up. Absolutely.

AUDIENCE: How?

JONATHAN GRUBER: I'll tell you that in a minute. But this is if the price goes up. OK. Other questions? It's a good question. Good clarification. OK. Now Let's go to the second effect, the income effect. The income effect is measuring the change-- the effect of the delta quantity due to income effectively rising. What do you mean income rises? Your income hasn't changed here. Why has income effectively risen?

Why are you effectively rich? I'm sorry. I'm sorry. The income has fallen. My bad. Your income is effectively falling in this case. Why are you effectively poor now? Your income hasn't changed, why are you effectively poor? Yeah.

AUDIENCE: Opportunity set is constricted.

JONATHAN GRUBER: Your opportunity set has been constricted. So your income itself hasn't fallen, but your resources have, what you can buy with that income has fallen. You're effectively poor. So it's not really that why has fallen, it's why over the composite price. The price level has fallen.

Why has the price levels fallen? The overall price level in society has increased. Pizza is the same, cookies are up. So your effective income is falling. Your real income has fallen. OK, you're effectively poor. What does that mean? If you're effectively poor, what's going to happen?

Well, what does that depend on? That depends on this. That's why I bother telling you so. I'm a big Doctor Seuss fan. Basically, if it's a normal good and your income goes down, what will you do? You'll have less of it.

But if it's an inferior good and income goes down, what will you do? You'll have more of it. So while the substitution effect can be signed, the income effect cannot. The effect cannot be signed because it depends on whether it's a normal or inferior good.

In the case we're showing here in figure 4.5, this is a normal good. So the income effect works in tandem with the substitution effect. The substitution effect, you want-- so you want fewer cookies for two reasons. Here's the intuition. Here's the key intuition. You want fewer cookies for two reasons.

First, cookies are a worse deal. Yet the opportunity cost of cookies has risen. Second, you're effectively poor. Your opportunity set is constricted, and cookies are a normal good. So for those two reasons, you want fewer cookies.

But what if cookies were an inferior good? Then the graph would look like figure 4.6. The graph will look like figure 4.6. In this case, the substitution effect would go one way, but the income effect would go the other way. And the net would be unclear.

Now we've drawn figure 4.6 such that the demand for cookies still falls. But you don't have to draw it that way. So the key thing is the substitution effect has to be the opposite of the sign of the price change to negative. But let's be clear, it's good clarification.

Substitution effect has to be the opposite of the sign of the price change. The income effect does not. The income effect depends on whether it's a normal or inferior good. This is pretty confusing. So I find it's useful to actually think about--

Well, let's actually rush through figure 4.6. Let's talk about what figure 4.6 is because it's not the-- I rushed through figure 4.6. Let's go back. Figure 4.6 is no longer pizza and cookies. Figure 4.6 is now steak and potatoes. Because pizza and cookies are both normal goods.

Potatoes is the classic inferior good. Right? Poor people, especially back in the day, used to have to eat potatoes because they were cheap to grow and very filling. And steak was unaffordable. But basically as you get richer, you want fewer potatoes and more steak Especially in agricultural society.

So let's say originally people are at BC1 where the price of steak is $5, the price of potatoes is $1. The price of steak is $5, the price of potatoes is $1. And let's say their income is $25. So in this example, income is 25. The price of a steak is 5 and the price of potatoes is 1.

So the budget constraint looks like BC1. You can do that for $25, you can have 25 potatoes or five steaks. And let's say your preferences are such that you end up at point A. You choose to have 3 and 1/2 steaks and 7 and 1/2 potatoes.

Once again, I just got that through the magic we showed you last time. This is a constrained optimization. I'm not writing down the utility function. I'm just going right to the results. But you know how to get there.

Now let's say that things change and the price of potatoes rise to $3 a pound. Steak is still $5 a pound, potatoes are now $3 a pound. Now, what happens is your budget constraint shifts into BC2. Your budget constraint is pivoted in massively. So now you still have five steaks and all you want is steak, but you can only have 8.3 potatoes on your $25 of income. Your budget constraint is BC2.

What happens? Well, let's look at the income and substitution effects. The substitution effect we get by drawing an imaginary budget constraint, BC prime. What are the features of BC prime? It is parallel to the new budget constraint, but tangent to the old indifference curve. And so that puts us at point B.

So substitution effect lowers our demand for potatoes from 7 and 1/2 to 4. We knew it had to lower it. Substitution effect is unambiguously signed. The price went up, so you got to have fewer of them. That's unambiguous.

The income effect, however, goes the other way because potatoes are inferior. In other words, now that you're so poor, now potatoes are so expensive, you don't have enough money for steak anymore. So you got to have more potatoes. It's ironic, right? But basic potatoes getting expensive has made you poor. When you're poor, you got to have more potatoes.

And so the income effect goes the other way because potatoes are an inferior good. So you, in effect, takes you back from wanting four potatoes up to wanting five potatoes. So the net effect, if we didn't talk about income substitution, would we just see--

We draw demand curve, we just say, oh, well, demand curves slope down. We get a higher price leads to fewer potatoes, but we'd miss the fact that the substitution effects are going opposite directions. Questions about that example?

AUDIENCE: So the potatoes, that's their price going up, right?

JONATHAN GRUBER: Yeah, they're all price going-- there's the price of potatoes going up to 3. I'm sorry, the price of potatoes went from 1 to 3. OK. Price went from 1 to 3. That led to a substitution effect that we know which way it goes. When a price goes up, you want less of something for the substitution effect. But an income effect, that went the other way.

Now that leads to the natural question, can you actually get an income effect for an inferior good that ends up bigger than the substitution effect? That is, could you actually get demand curves to not slope down?

Think about what would happen if the income effect was so big that it pushed you past the original point? It pushed you to the right past the original intersection? Well, that would imply something funky. That would imply demand curves slope down. Demand curves slope up instead of sloping down. That imply demand curves slope up.

That would say higher prices make you want more of something, which violates what we-- the very first lecture of this class, we said demand curves slope down. But it's possible theoretically that you can get what's called a Giffen good.

A Giffen is a good where the income effect, the good is so inferior that the income effect dominates the substitution effect you end up with a wrong signed relationship between price and demand. You actually get that a higher price leads to more demand, or an upward sloping demand curve.

Now, can that happen? Well, I like the name Giffen because it rhymes with Griffin. And Griffins are imaginary. Giffen goods are also largely imaginary. It's a great theoretical concept, but they don't really exist. Yeah.

AUDIENCE: A Veblen good.

JONATHAN GRUBER: A Veblen goods is a different-- Veblen, that's a totally different concept. That's the concept that basically you want more when the price is up because you care about your social standing. But that's a behavioral economics thing. We'll come back to that type of behavior. It's not a standard economics effect.

So Giffen goods really don't exist. The classic example people often had was, well, gee, when the price of potatoes went up in Ireland because of the potato blight, people ended up consuming more potatoes. So potato consumption went up per person in Ireland when the price of potatoes went up during the potato blight.

People said, aha, potatoes are a Giffen good. But here's what they got wrong. Here's where empirical economics is fun. They miss the fact that what other effect did the potato blight have? It chased all the people who could afford it out of Ireland. The only people who were left were super poor. So, yeah.

So yeah, on average, potatoes per person went up. But it wasn't that each person consumed more. It's just the people who didn't consume many potatoes left. And all the people left were just huge potato consumers.

So if you look at the data, you see potatoes per capita went up. It's a Giffen good. What you're missing is the capitas are different. Each person needs more potatoes. It's just the non-potato eaters left. Those people left were the potato eaters. And that's why empirical economics is fun because it can help explain these kind of anomalies.

So where does this leave us? Let's draw a chart here to explain this all. So this, let's imagine we're going to have here the price change. The price change. We're going to have the substitution effect, the income effect, and the net effect.

And we're going to start with normal goods. We start with normal goods. So with normal goods if the price goes up, like I just showed you, the substitution effect must be negative. It must be the opposite of the price change. The income effect, if it's a normal good, must also be negative. Therefore, on net, you get a higher price leads to lower demand.

Likewise, if the price goes down, this is what this good question was highlighting. I was just looking at the price going up. Well, in that case, the substitution effect-- let's think about the intuition. If the price of cookies went down, we have just lowered the opportunity cost of having cookies. Therefore, we're going to want more cookies and fewer pizza through the substitution effect.

So therefore, the substitution effect will be positive if the price goes down. The income effect will also be positive. Why? Why would the income effect will be positive if the price of cookies goes down? Why would the income effect be positive? Yeah.

AUDIENCE: Because you're richer.

JONATHAN GRUBER: You're effectively richer. You can buy more cookies with the same money. Since you're richer, you want more normal goods. Cookies are normal good, you want more cookies. So this goes up.

Now what's more interesting is what if these goods were inferior? What if cookies were inferior? So now when the price goes up, what happens to the substitution effect? The substitution effect, it goes down, as this gentleman helpfully pointed. It goes down.

Why? Because we know the sign of substitution effects are always the opposite of price change. No mystery there. Price goes up, substitution effect drives demand down. Price goes down, substitution effect-- I'm sorry, goes up. Yeah, right.

Price goes up, substitution effect drives demand down. Price goes down, substitution effect drives demand up. Those never change. So here, if the price goes down, substitution effect would be positive. Those don't change.

What changes is the income effect. If it's an inferior good, if it's an inferior good, now when the price goes up, when the price goes up, you want more of it. And when the price goes down, you want less of it. So the net is unclear.

So once again, if goods are normal, then demand curves look like-- then basically there's no-- if goods are normal, it seems like I just wasted 20 minutes of your time. Right? because if goods are normal, it's like, well, look, yeah, they build on each other, but we end up in the same place and the constraints are imaginary anyway, so who cares?

The answer is, I'll tell you why we care in three or four-- in about six lectures. Because it turns out when you make your decisions about how hard you're going to work or how much you're going to save, income substitution effects are going to start to work against each other.

This case is only really interesting if the income effects-- if it's an inferior good, and then you get some interesting things. Although in practice, we don't really have Giffen goods. But in theory, at least you can get interesting things. OK? Questions about that?

OK. One other announcement-- please don't go. I really love that you guys come after lecture and ask me questions. That's awesome. I just do have to say on Mondays, I have to book out of here fairly quickly, so don't feel dissed. I do have office hours. And if you can't make the office hours, just email me. I'm in the office lots of times. So just email me. We can talk. Thanks.

AUDIENCE: Thank you.